Sandia National Laboratories

Fact Sheet

High-Performance Computing

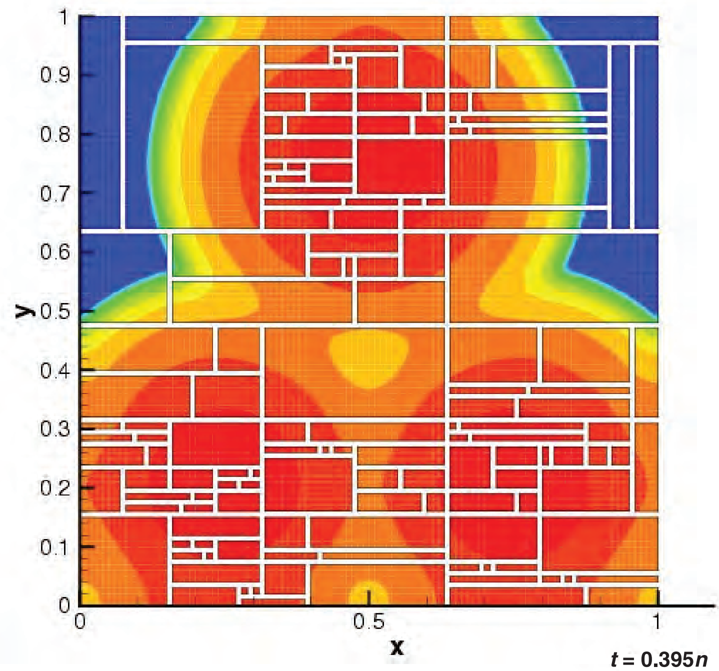

Figure 1. Temperature concentration profile of an igniting H_2–air mixture. Simulation was performed on an adaptively refined mesh using the CFRFS toolkit, which is based on the common component architecture.

Within the scientific community, high-performance algorithms and parallel computing are growing in significance because of the increases in demand for numerical computation and in the availability of massively parallel computational clusters. Scientists in a growing range of fields can now use large computational clusters to model and solve complex and computationally intensive problems. The change from symmetric multiprocessing systems to multinode distributed systems has brought many opportunities and challenges. Scalable computing research and development (SCRD) researchers have expertise ranging from low-level hardware to advanced scientific- algorithm development. Our researchers have made significant advances in addressing these challenges by leading the path to ultrafast InfiniBand interconnects, designing scalable algorithms while managing software complexity, developing software tools to measure CPU and communication performance, and collaborating with the high-performance computing (HPC) community on the development of critical components in the software stack. We are also on the forefront of R&D into high-speed scalable storage systems to meet the needs of terascale data-intensive applications.

Scalable Algorithms

Employing an object-oriented design and parallel-programming techniques, developers of the massively parallel quantum chemistry (MPQC) package are working to mitigate the costs of implementing and running large scientific applications. These researchers are on the forefront of scalable quantum-chemistry algorithm design. Recently, Sandia researchers developed massively parallel algorithms for reduced-scaling correlation methods, which circumvent the nonphysical, higher-order polynomial scaling of conventional methods. As the only software package that can perform such computations on massively parallel computers, MPQC provides the capability to compute, with high accuracy, energies for molecular systems of unprecedented size.

Common Component Architecture

The common component architecture (CCA) provides a means for software developers to manage the complexity of large-scale scientific simulations and move toward a plug-and-play environment for HPC. Members of the SCRD department at Sandia National Laboratories are key participants in CCA research, which spans from standards and framework development to adoption within application areas such as quantum chemistry.

SAND2006-2985P
April 2006

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Reconfigurable Logic

The availability of inexpensive, reconfigurable computing signals a paradigm shift in how data-intensive calculations can be carried out faster and with much lower power requirements than can be achieved with traditional, general-purpose computing. We have been evaluating the use of reconfigurable, field-programmable gate arrays (FPGAs) to build specialized hardware for certain types of scientific simulations. Our research interests include the use of FPGAs to perform fast grid interpolations in particle simulations. Grid scoring is used extensively in various forms of ligand design, computational chemistry, solvation modeling, and smooth-particle hydrodynamics. In these simulations, streaming tens or hundreds of millions of particle positions through the grid evaluator is common, resulting in very significant amounts of computer time. On an FPGA, these evaluations can be pipelined almost perfectly, resulting in very high throughput on the FPGA and permitting the dedication of the general-purpose CPU to optimization and particle–particle interactions.

Performance Tools

One downside to using distributed supercomputers is the difficulty of designing and programming efficient algorithms. An application's performance is highly dependent upon how a programmer organizes computationally intensive algorithms and uses message-passing libraries for concurrent execution among many processors. To assist developers who use distributed computer systems, SCRD researchers are developing software tools for evaluating and enhancing application performance. Specifically, these Sandia researchers are focusing on hardware-performance monitoring and communication-performance analysis.

InfiniBand

Simulation work at Sandia National Laboratories demands powerful but economic computing capabilities. Researchers at Sandia have been working to address these needs by exploring advanced architectures for cluster computing. InfiniBand, a commodity multivendor open-standard network interconnect, has been proven to have the desired characteristics for a high-performance cluster interconnect: high bandwidth (10 to 30 Gbps currently), low latency (1 to 3 µs), low CPU overhead, and remote-direct memory access (RDMA). Our current research focuses on exploiting InfiniBand's advanced communication features in Sandia's clusters.

OpenIB

SCRD researchers are key participants in the development of a common multivendor InfiniBand software stack, known as OpenIB, that emphasizes high performance, scalability, and interoperability, while reducing deployment time. By promoting an open-development model, the OpenIB alliance continues to encourage community participation. In addition, OpenIB is in the process of becoming the standard RDMA device stack in the Linux kernel.

Open MPI

Open MPI is an open-source, freely available implementation of the MPI-1 and MPI-2 interprocess communication standards. Open MPI's design goals include high performance, fault tolerance, multiple concurrent network support, and ease of use. SCRD researchers are collaborating in Open MPI development and are currently investigating InfiniBand and Red Storm support for Open MPI.

Scalable Storage

Storage platforms have failed to keep pace with advances in computing power. Today's high-speed clusters can easily use the latest interconnect technologies for node-to-node communications, but all too often, the input/output storage systems are still bottlenecks. A versatile, scalable, sharable, and secure storage architecture is needed to balance the computational infrastructure for HPC. SCRD researchers are exploring parallel-access techniques and RDMA to improve scalability. Sandia scientists are also promoting a standardized common file-system interface to facilitate sharing, while ensuring compatibility with Sandia's corporate security infrastructure.

Multiscale Simulations

In recent years, multiscale/multiphysics simulations have emerged as one of the most viable candidates for extending simulation capabilities to handle fine-grained events in simulation systems of physically relevant size- and time scales. Sandia researchers have been at the forefront of creating methods and software for multiscale/multiphysics simulations, and they continue to collaborate with other research groups to explore applications of these new tools.

For more information, contact
Scott Cranford, (925) 294-3617